MULTIMEDIA UNIVERSITY

STUDENT ID NO

MULTIMEDIA UNIVERSITY

FINAL EXAMINATION

TRIMESTER 1, 2017/2018

TSN2101 – OPERATING SYSTEMS

(All sections / Groups)

13 OCTOBER 2017
9.00 a.m - 11.00 a.m
(2 Hours)

INSTRUCTIONS TO STUDENTS

1. This Question paper consists of 8 pages with 7 Questions only.

2. Attempt **FIVE** out of **SEVEN** questions. All questions carry equal marks and the distribution of the marks for each question is given.

3. Please print all your answers in the Answer Booklet provided.

QUESTION 1

a) Co-operating processes can affect or be affected by other processes, including the sharing of data.

i) What are the FOUR reasons for providing an environment in operating systems that allows for co-operating processes?
ii) What are the two fundamental models of inter-process communication (IPC)?

[2 + 1 = 3 marks]

b) Define a process and describe what is meant by short-term and long-term scheduling of processes.

[3 marks]

c) In the development and design of most modern operating systems, a modular kernel may be the best of the design techniques. Explain why this is so.

[4 marks]

d) Operating systems (OS) are interrupt driven. How interrupts are handled by the OS?

[2 marks]

Continued ...

QUESTION 2

a) What is the difference between preemptive and non-preemptive CPU scheduling?
[1 + 1 = 2 marks]

b) Consider the following set of processes:

Process	Arrival time	Burst-time	Priority
P1	0	9	5
P2	1	4	2
P3	2	3	1
P4	4	1	3
P5	5	6	4

i) Draw the Gantt charts illustrating the execution of these processes using **Preemptive Priority** (a smaller priority number implies a higher priority), **Non-Preemptive Shortest Job First (Non-Preemptive SJF)** and **Round Robin (RR) (quantum = 3)** scheduling.
[2 + 2 + 2 = 6 marks]

ii) What are average turnaround times for **Preemptive Priority and Non-Preemptive Shortest Job First (Non-Preemptive SJF)** algorithms?
[2 + 2 = 4 marks]

Continued ...

QUESTION 3

a) An object 'mutex' of the class Semaphore is used to achieve mutual exclusion for the critical section of the codes. Functions of the class include `acquire` (`mutex.acquire()`) and `release` (`mutex.release()`), corresponding to acquiring and releasing the locks in critical sections of codes respectively. Show where these functions are used in the codes below:

```
Semaphore mutex;
Do{
      \\critical section
            ⋮
      \\remainder section
} while (TRUE)
```

[3 marks]

b) Consider the following system with **FIVE** processes and **THREE** resource types:

- A has 12 instances
- B has 12 instances
- C has 12 instances

The snapshot of the system is given as:

Process	Allocation A B C	Max A B C
P0	0 2 2	10 10 12
P1	2 2 2	6 4 6
P2	2 2 2	8 6 8
P3	2 2 2	4 2 4
P4	4 4 2	10 8 10

Available A B C
2 0 2

Answer the following questions using the banker's algorithm:

i) What is the content of the matrix *Need*?

[2.5 marks]

ii) Find a sequence of processes that will get the system into a safe state, if there is any. Show all required steps.

[5 marks]

iii) P4 makes a request for 2 instance of A and 1 instance of C. Check whether the request can be granted.

[1.5 marks]

Continued ...

QUESTION 4

a) In paging, pages in memory are tracked through the use of page tables. These page tables are also kept in memory.

i) Briefly state how these page tables are accessed by the OS with the use of page table base registers (PTBR) and page table length register (PTLR).
ii) Since page tables are also kept in memory, every data or instruction access requires two memory accesses. Briefly discuss how this problem can be solved.

[2 + 2 = 4 marks]

b) Assume that a system has a TLB hit ratio of 90 %. It requires 15 ns to access the TLB, and 85 ns to access main memory. What is the effective memory access time in nanoseconds for this system?

[2 marks]

c) Consider a computer system with a 24-bit logical address and 4 KB (1 KB = 1024 = 2^{10}) page size. The system supports up to 8192 MB of physical memory. How many entries are there in each of the following:

- A conventional single-level page table
- An inverted page table

[1 + 1 = 2 marks]

d) Identify the page numbers and offsets for the following address references (Given as decimal numbers). Assume the page size is 2 KB (1 KB = 1024).

- 20466
- 40000
- 512
- 6264

[0.5 + 0.5 + 0.5 + 0.5 = 2 marks]

e) Consider a system that uses a three-level page table with page size of 1 KB and 48-bit logical addresses. The first 12 bits of the address serve as the index into the first-level page table, and the last 12 bits serve as the index into the third-level page table.

i) How many bits are needed to specify the second-level index?
ii) How many entries are in a level-one page table?
iii) How many entries are in a level-two page table?
iv) How many pages are in the virtual address space?

[0.5 + 0.5 + 0.5 + 0.5 = 2 marks]

Continued ...

QUESTION 5

a) Virtual memory is a memory management capability of an OS that uses hardware and software to allow a computer to compensate for physical memory shortages by temporarily transferring data from random access memory (RAM) to disk storage. Describe any two advantages and two disadvantages of implementing virtual memory systems?

[2 + 2 = 4 marks]

b) Consider the following page table in memory.

Frame	Valid/Invalid Bit
5	i
2	v
4	v
20	i
15	i
8	i
9	v
11	i

Assume that the page size is 512 Bytes, check whether the virtual addresses below generate page faults. If the address does not generate a page fault, identify its physical address.

i) 656
ii) 1444
iii) 2570
iv) 3604

[2 marks]

Continued ...

c) Consider the following page reference string:

5 2 3 4 2 5 6 2 3 4

Assume the availability of **FOUR** frames and all the frames are initially empty. Consider the **First-In-First Out (FIFO)**, **Least Recently Used (LRU)** and **Optimal (OPT)** page replacement algorithms for Virtual-memory management.

- Identify the number of page faults that would occur if FIFO is used.
- Identify the number of page faults that would occur if LRU is used.
- Identify the number of page faults that would occur if OPT is used.

Show all steps involved.

[2 + 2 + 2 = 6 marks]

Continued ...

QUESTION 6

a) List down two benefits of contiguous disk block allocation for files in a file system.

[2 marks]

b) What are the possible implementation problems in contiguous disk block allocation?

[2 marks]

c) How fragmentation can improve the performance of a file system?

[2 marks]

d) What is the difference between internal fragmentation and external fragmentation?

[2 marks]

e) A 4 TB hard disk with a 4 KB block size is using bitmap to maintain free space list. How much disk space is needed to store the bitmap?

[2 marks]

f) In linked allocation, each file is a linked list of disk blocks: blocks may be scattered anywhere on the disk. Explain why linked allocation does not perform well in random file access.

[2 marks]

Continued ...

QUESTION 7

a) Suppose that a disk drive has 2000 cylinders, numbered 0 to 1999. The drive is currently serving a request at cylinder 1145, and the previous request was at cylinder 1125. The queue of pending requests, in FIFO order, is

60, 1455, 820, 1760, 935, 1410, 1110, 1815, 120

Starting from the current head position, what is the **schedule** and the **total distance** that the disk arm moves to satisfy all the pending requests, for each of the following disk-scheduling algorithms?

i. LOOK
ii. C-SCAN

What are the **factors** influencing the selection of a disk-scheduling algorithm?

[3 + 3 + 3 = 9 marks]

b) Windows schedules threads using a priority-based, preemptive scheduling algorithm. Discuss how Windows' dispatcher perform this.

[3 marks]

End of Exam